If Linux Were Just Getting Started Today, It Would Ge Crushed, and We'd All be a Lot Poorer for It. We Can't Let that Happen to AI.

If Linux were just starting out today, it would get crushed.

It would get killed by lawmakers, skeptics, doomers, alarmists and vast array of hysterical conspiracy theorists spinning terrifying fantasies of what happens if we let software be open and free.

They'd scream that open source was enabling our enemies and it's too dangerous to let into the hands of regular people and it's got to be tightly controlled or *really* bad things will happen! They'd spend billions and fire up idiotic law after idiotic law at the state and federal level to make Linux kernel developers responsible for "harms" aka whatever nebulous, ill-defined, anti-democratic, anti-choice thing they could imagine.

And the entire world would be so much poorer and worse for it.

Why? Because Linux and open source is the invisible backbone of the digital world.

It powers 100% of the top supercomputers in the world, the clouds of Amazon, Microsoft, and Google, US army laptops, many US Navy sonar systems, the router in your house, your smart phone, Ham radios, the Playstation 5 (powered by FreeBSD, another open source OS), and so much more.

A report by a team at Harvard Business School estimated the value of open-source software to be around **$8.8 trillion annually to the global economy**. Rebuilding Linux's codebase and infrastructure would require **$4.15 trillion** (supply-side) to **$8.8 trillion** (demand-side) in investment.

But if Linux were starting today it would get murdered before it even crawled out of its crib. Trillions of dollars in economic value would go right up in smoke.

And that's because today, there's just too many hostile forces arrayed against openness.

If you tell today's politician that we'd get those trillions in economic value plus the strategic hard and soft power advantages of having the entire world run on our software but China could use it too, they'd say no, we we've got to stop it, not understanding that there is simply *no way to get all that without having it open in the first place*.

Open is the very essence of what drives all this value and there is no way to get it with partially open, or kind of open, or sort of open, or just outright closed.

How do we know? Because we tried semi-open in the early 90s with encryption and it simply didn't work.

US encryption export controls in the early 1990s crushed early e-commerce development by forcing companies to use weakened security measures. That's because short-sighted Cold War-era policies classified encryption as munitions, and required export licenses for cryptographic tools exceeding 40-bit key lengths. Browser makers like Netscape had to release a 128 bit encryption version for Americans and a weaker 40 bit encryption version for everyone else.

SCO EXAMINER ★★★ Tuesday, September 21, 1993 D-3

Grand jury is investigating encryption software exports

ASSOCIATED PRESS

SAN JOSE — A federal grand jury is investigating exports of a controversial computer program in a case that could affect how software is distributed worldwide.

U.S. Customs officials asked for an investigation into ViaCrypt of Phoenix and Austin Code Works of Austin, Texas, and the companies' plans for foreign distribution of software, including PGP, a program that turns data into an indecipherable code using encryption technology.

William Keane, an assistant U.S. attorney, confirmed that an investigation is continuing, but declined to comment on the case.

The PGP program has been distributed worldwide over computer networks by some computer enthusiasts who oppose the U.S. government's trade regulations on encryption.

The National Security Agency, which monitors international communications, has supported strict encryption technology export regulations, arguing that it would be difficult to keep tabs on hostile governments and foreign terrorists.

But opponents say the restrictions hurt sales and violate the First Amendment that protects the right to publish information about encryption.

"I wrote PGP to make democracy healthier. I didn't do it to make money," said Philip Zimmermann, a computer consultant who developed PGP.

No trial for exporter of privacy program

U.S. relents on writer of 'dangerous' encryption software

FROM EXAMINER STAFF AND WIRE REPORTS

his role in writing and popularizing Pretty Good Privacy. The program, which is available free over the Internet, lets computer users turn their e-mail into a jumble of numbers and letters unreadable to anyone except the intended recipient.

The PGP code is so unbreakable that it has been classified as a "dangerous munition" under the International Traffic in Arms Regulations in the Arms Export Control Act. The ITAR makes it a federal crime to export munitions

[See PRIVACY, B-2]

(Source: San Francisco Examiner)

Unfortunately, that just ended up compromising American security too. Why? Because the hoops you had to jump through to get the high encryption browser meant that many Americans didn't bother and just used the weaker version too, putting millions of people at risk with easily and laughably hackable "security." These rules weren't reversed until the late 1990s.

Until we had strong encryption worldwide e-commerce simply couldn't take off because who would trust their credit card to something insecure? Outside of browsers, it also cost America billions in foreign sales to our allies and friends in places like the EU and the UK. Banks and other security demanding industries overseas rightly demanded better encryption and so turned to local software makers who could offer them the security they needed to run their business right.

Today, global e-commerce is between **$6.56-$7.4 trillion** annually, demonstrating exponential growth despite these early obstacles.

How much more would it have grown and how much faster if we'd taken the chains off much earlier?

This is what short-sighted, alarmist, jingoistic policies do. They hold back the whole world and they blow back on the country making them too.

Luckily when Linux was getting started, it was a different world. It was a world where software was free to develop naturally and organically without much interference from a droning chorus of activists and antagonists because few people understood technology and there was no social media to amplify alarmist's voices.

Beyond that, the US had just won the cold war. It was the one true power on Earth with no realistic rivals. China was still a poor backwater, crawling its way out of the disastrous Mao era and ruinous communist social experiments like the Great Leap Forward and the Cultural Revolution. China's GDP in 1991 (the year Linux was released) was $383.37B. By contrast, the US had 6.158 trillion in GDP in 1991 and China had 18.80 trillion GDP in 2024 (with the US at 29 trillion in 2024).

The first Gulf War shocked militaries around the world and everyone feared the awesome might of the US military which rolled over the 4th largest army in the world like it was a video game in a few months.

(General "Stormin'" Norman Schwarzkopf, architect of the first Gulf War)

The software tech industry was small and the average person didn't have any opinions about it, much less any understanding of the web. Politicians like Clinton embraced tech as an engine of growth and prosperity (with notable exceptions like the v-chip). Lawmakers knew they didn't know much about technology and generally stayed out of it, letting it develop naturally and organically, with a few minor exceptions. No one would pass Section 230 today. Despite it having made the modern web what it is today it's hated on both sides of the isle in American power politics.

Today we have a much, much different world. It's a multi-polar world, with China ascendant and rivaling or surpassing American economic power in multiple areas. China's car industry went from non-existent to the number one car manufacturing engine in the world in just twenty years, producing 30 million vehicles a year, double what the US produces and at 20%-25% of the cost of the rest of the planet. DeepSeek shocked the world and showed many Americans that China wasn't behind on AI at all and may even lead it soon, if we're not careful.

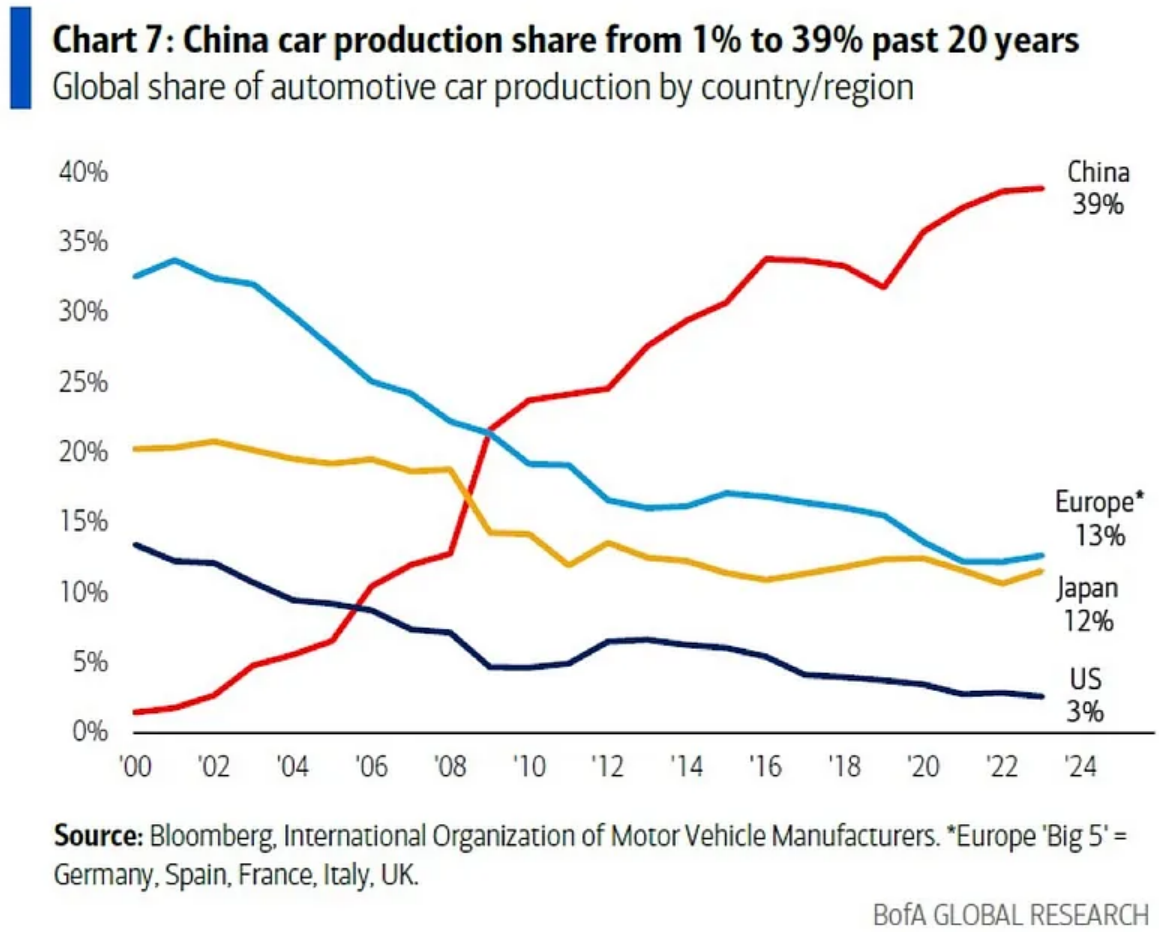

(Source: BofA Global Research)

But the biggest difference today is that we have a whole range of people, from lawmakers, to social pundits, to the local cab driver with zero understanding of software, AI, and technology and who take incredibly aggressive stances on it.

Hundreds of bills are pressing forward at the state level, creating a tangle of liability, paperwork and rent seeking compliance officers who produce no economic value by driving up costs by 5-15% across the board. We're averaging 19 new bills *per day* and looking at over 1000 bills by the end of 2025. Make no mistake, the EU AI Act is Coming to America.

And that's because for over two years, this band of doom-sayers and alarmist have hijacked public discourse—spinning sci-fi doomsday fantasies fueled by personal insecurities, convinced that they'd unearthed humanity's ultimate threat and they were the only ones who could see clearly and stop it!

ChatGPT was a catalytic moment for Doomers and alarmists of all stripes. They seized the moment as ChatGPT's exploded onto the world stage to shout that AI was about to outsmart us all and we needed to act before it was too late!

The alarm spread like wild fire and they fanned the flames of a full blown moral panic.

Unfortunately, too many people, especially lawmakers, listened and the fallout was a bunch of clumsy, shortsighted policies like the EU AI Act and boondoggle events like UK's AI Safety Summit, a virtual circus of grifters and charlatans taking about utter nonsense like AI snuffing out all life by causing solar flares or sucking all the oxygen out of the world by gene editing plants. I wish I were making both of those up.

Europe suffered the worst because it listened to some of this early insanity. They pushed out Thierry Breton, the architect of the EU AI Act, too late. The damage was done. Now Europe faces the looming arrival of the law and the overwhelming burden of outdated, restrictive rules that stifle innovation without addressing any real world, practical risks like autonomous weapons or mass scale surveillance.

But who are these ill-informed and highly aggressive anti-open source/weights, anti-AI alarmists and how do we defend against them? Broadly, these folks fall into one of four categories:

- **The monopolist**
- **The jingoist**
- **The doomer**
- **The idiot**

One could argue that the last category is really just a variation of the first three but the idiot deserves his own category since idiots produce a unique kind of argument that must be defended against if we're going to defend open source in general.

Future History is a reader-supported publication. To receive new posts and support my work, consider becoming a free or paid subscriber.

Monopoly Money

Let's start with the monopolist.

This is the most straightforward antagonist of open-source and open weights AI. They have a simple and clear motivation: money.

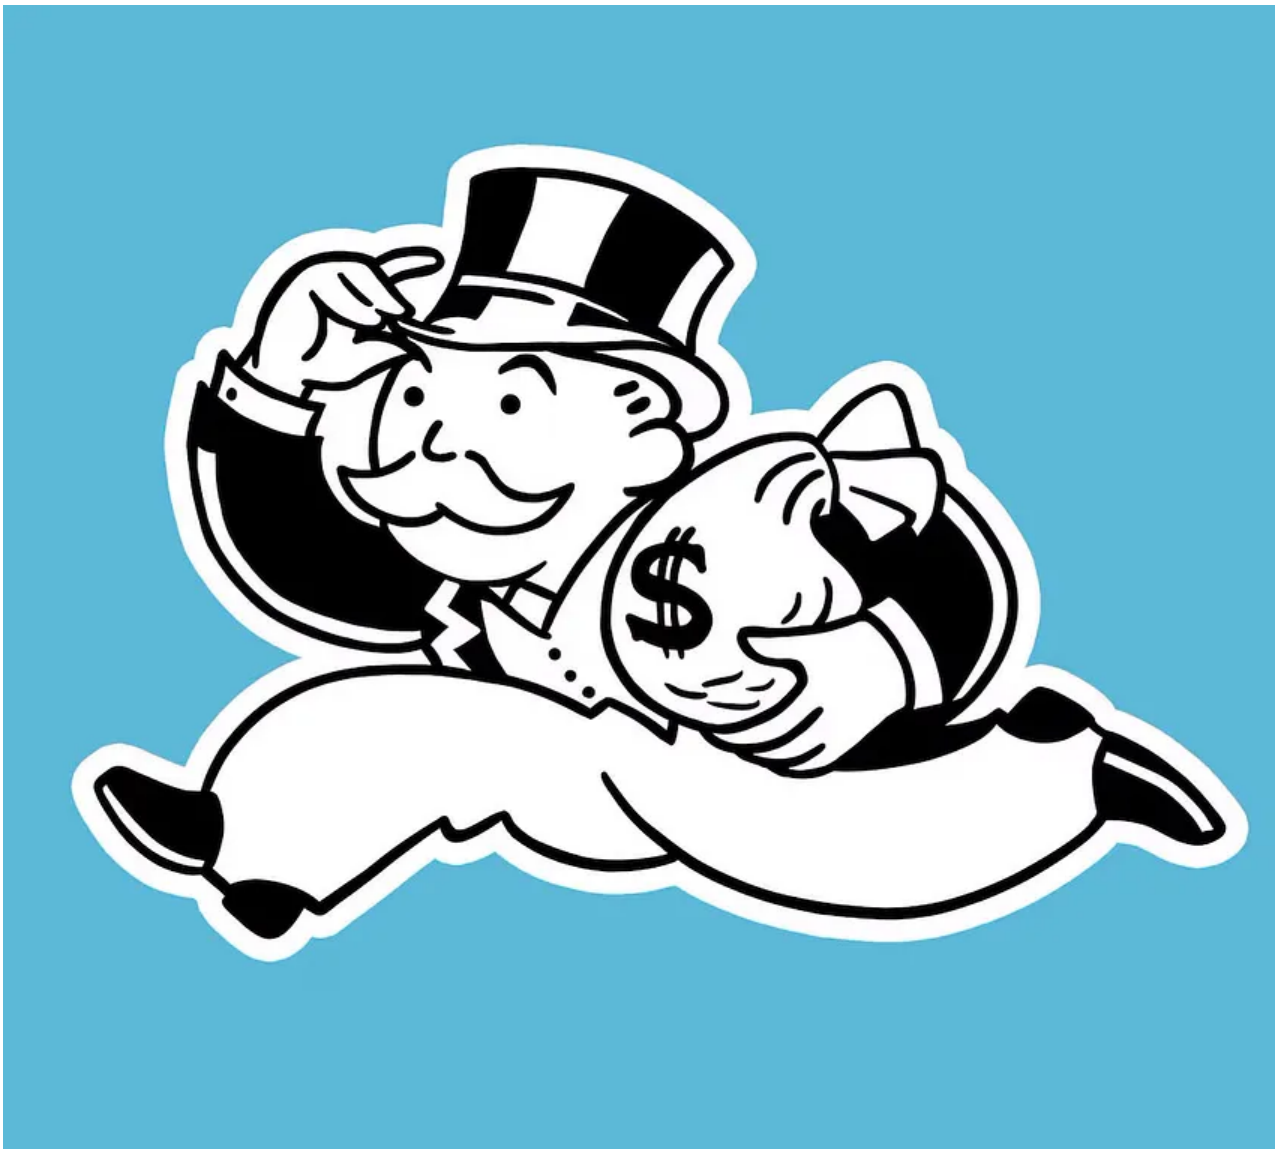

Now don't get me wrong, there's nothing wrong with money. I like it and so do you and if you're say you don't you're lying. But the problems start when the monopolist doesn't want to play fair with competition. They'd rather ensure the competition doesn't exist or gets locked out of the marketplace via the law so they can't compete at all.

Monopolists will often pretend to like open source or pay lip service to it, but in general they despise it for one major reason:

Open source tends to drive down prices.

Monopolists would love to keep the prices high at all costs because they have a tremendous investment in R&D and a large install base.

Sam Altman who's taken a decidedly closed source approach with OpenAI wrote in a Reddit AMA, "I personally think we have been on the wrong side of history here and need to figure out a different open source strategy. Not everyone at OpenAI shares this view, and it's also not our current highest priority."

In other words, lip service. Open source is great but we're not really going to make any tangible shift towards it because it's not a current priority and it would likely gut their current pricing structure over time.

OpenAI's biggest rival, Anthropic, is more openly monopolistic. Half of CEO Dario Amodei's recent essay on DeepSeek and China was about prices and how we have to act fast to make sure China doesn't become a real AI superpower so we need deep sanctions and chip controls to keep them out of the market. But his main concern was that cheaper, open models would be good enough, less controlled and less opinionated than Anthropic's models, which would slice into their bottomline while dressing it up an national security.

The other big reason monopolists hate open source is because open source tends to make software a commodity. It makes it widespread and ubiquitous. It's permissionless and doesn't require anyone to go through centralized choke points to get it. It lets anyone dream up any way imaginable to use it and just do it without begging for the right to do it. Which again, means it comes down to price. It's much harder to charge top tier prices when everyone has a common set of building blocks that makes it easier to move "up the stack" because the pieces below have become commodity prices, aka cheap.

DeepSeek is already eating into OpenAI's bottom line. R1 was finedtuned by Perplexity to strip out the model's unwillingness to answer questions about sensitive Chinese history, like Tiananmen Square, (due to compliance with local Chinese laws) and it almost certainly powers the $20 a month rival to OpenAI's $200 a month DeepResearch. DeepResearch is an amazing product and I'm a subscriber and lover of the platform but Perplexity's Deep Research clone, using a commodity intelligence layer, is likely good enough for many people's everyday needs, which means less people sign up for the more expensive subscription that is out of the range of most people's budget anyway.

Let's face it, the average person simply does not have $2,400 of spare cash a year to spend on AI apps.

Cheaper intelligence means more choice for consumers and more options. That's good for everyone except the monopolist.

But since the monopolist can't come out and say what they really want they have to frame their argument differently so they usually swing like Errol Flynn into the **Moral High Ground Fallacy**. This is a super simple attack and it works by appealing to vague moral principles or platitudes like "fairness" or "justice" or "compassion" and deflecting away the weakness of their argument by framing their opponents as unethical scum instead of addressing the substance of the argument.

When it comes to AI those two moral arguments are usually "safety" or "China."

Safety is a clever word. It's a universal platitude that has some real meaning but that mostly just allows anyone to project their own feelings and ideas into the word. It's like if someone says the word "God." People say the word and think everyone else is hearing

the same meaning they have but really what other people are hearing is something completely different. It's colored by their own ideas and faith or lack thereof and so the word immediately carries a distorting charge that warps its meaning in conversation.

The same is true for a word like safety. Who doesn't want to be safe? If I offer you a choice between extreme danger and sitting at home comfortably safe on your couch watching TV, you're going to chose safe every day.

But what is safe in AI? Is it simply the inability for the system to take dangerous actions? Is it the refusal to answer certain kinds of questions? Who defines those questions? If I am a writer and I want to know the chemical formula for a poison to make a killer character in my story more realistic, how will the AI know the difference? Should it be making those kinds of assessments anyway? Who gets to define the rules? How do we know they're right? It quickly becomes a quagmire and very soon you start to realize that safety and censorship have a whole hell of a lot in common. They might even be the same animal.

Safety in AI usually comes down to giant fantastical fears like AI launching nukes or bioweapons, the kind of nonsense that comes from Hollywood films and has nothing to do with reality. Any idiot who's asked "do you want AI to control nukes or launch a deadly virus" is going to say no, absolutely not. Who in the world would want such a thing? The main problem is that it's a nonsense question in the first place. AI can't do these things now. It's just imaginary made-up speculation that may be possible at some point in the future if someone was stupid enough to link up an AI control system to a nuke but that has absolutely no relevance to safety discussions today.

So you scare people with a vague future threat and get them to push for closed, tightly controlled systems, which plays right into the ready and waiting hands of the monopolists.

> "You don't want unsafe AI do you? Like what if AI killed you and everyone you loved? You wouldn't want that would you? Well we can keep you safe! All we have to do is ban openness, ban competition and we'll deploy special, trusted systems that only we, the trusted people, can control to make you safe."

Beware of monopolists bearing gifts.

The other vague trigger word for monopolists is "China".

China is the bugaboo of the left and right leaning defense hawks alike. It's today's Soviet Union, a big bad guy to pin everything on. Amodei's essay called for ramping up export controls on US AI chips from NVIDIA in order to create a "unipolar world" where the US and its allies hold all the cards in AI versus a multipolar world where China is an AI superpower too.

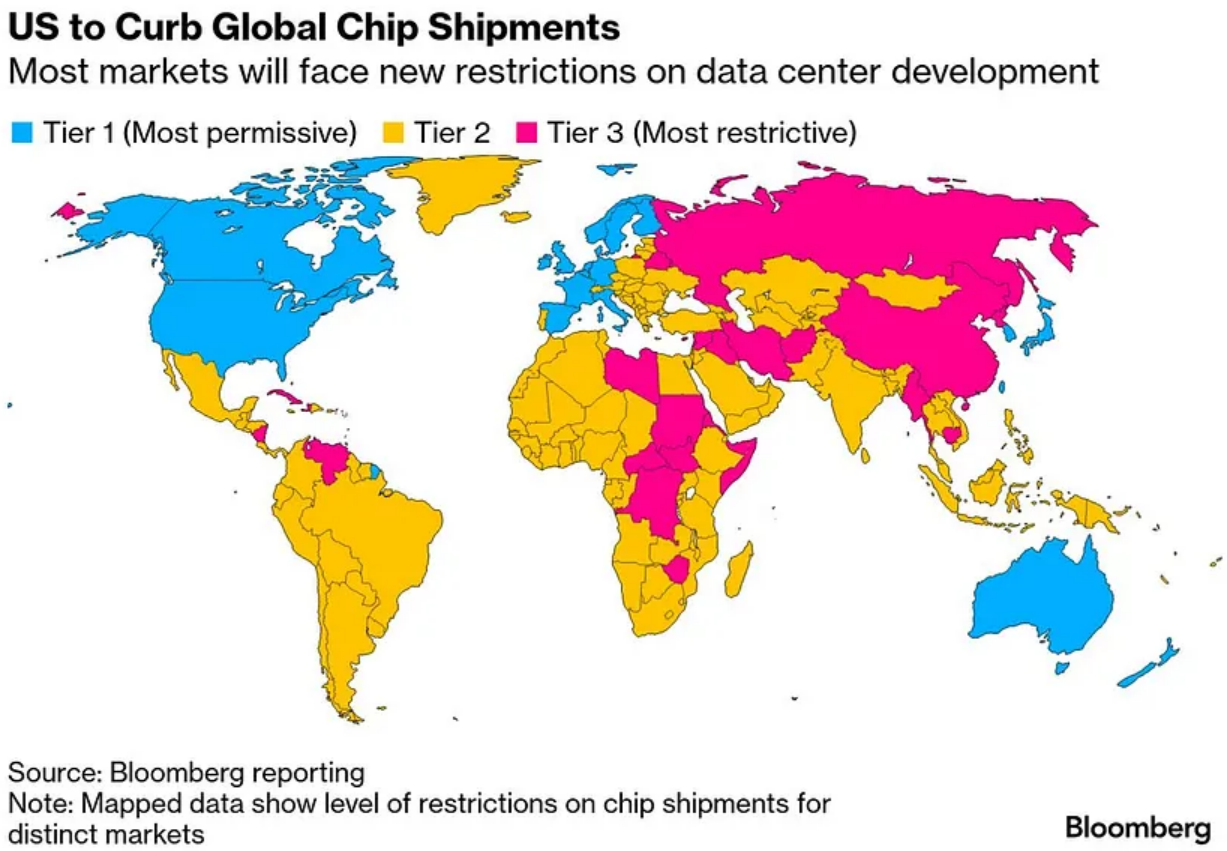
US to Curb Global Chip Shipments
Most markets will face new restrictions on data center development
Tier 1 (Most permissive)
Tier 2
Tier 3 (Most restrictive)
Source: Bloomberg reporting
Note: Mapped data show level of restrictions on chip shipments for distinct markets
Bloomberg

(Source: Bloomberg, Export Controls Map by Biden and Trump administrations)

AI card sanctions don't and won't work.

Chinese companies can easily set up shop in any country on Earth because today's businesses are multi-national by default. Set up in the US, Europe or Singapore and they can train models with NVIDIA chips and run them to their heart's content.

And Chinese companies are not sitting still. In a striking parallel to earlier encryption export controls, where American companies were locked out of competing in foreign markets, we've blocked our American champion, NVIDIA, from selling into the Chinese market with increasingly drastic export controls, and we've create a void that Chinese companies like Huawei are leaping in to fill. DeepSeek inference is already running on Huawei Ascend 910Cs and Huawei is already previewing 920C chips that aim to compete with Nvidia's upcoming Blackwell chips in terms of power and speed.

This is a losing game of tit-for-tat trade war, where we block them from getting chips, they make their own and then Beijing hits back and block us from key rare earths like Gallium that we need for everything from our special forces night vision goggles in our military to circuit boards.

It's a war of attrition with no winner, ever. We will have a multipolar world, whether people like it or not and the only thing that increasing the trade war will do is lead us to a monopoly of closed source American firms who leverage the China fear to get restrictions on competition.

Remember that ultimately the base goal of the monopolist is to reduce competition and control the market outright by having the government pass laws that favor their business and crush competitors. This creates a small group of a few companies who can do whatever they want, even if it's not in the best interests of small business or consumers.

American's business and technological prowess over the last few centuries came from small businesses that were able to operate and innovate as they saw fit with minor restrictions that made sense (with notable exceptions like AT&T in the 1950s which was a state sponsored monopoly). Lots of smaller companies compete and eventually some of them grow into massive companies. Kill that in favor of a few mega-companies and innovation slows down to a trickle as innovation gets tightly restricted to the limited imagination and gate keeping of one or two companies.

Already, AI has become a closed environment where the big American AI companies release no papers with any useful details and provide nothing back to the research community that they benefited so strongly from in the first place. Without kubernetes, Ray, containers, Linux, pytorch, the Transformers architecture, Claude and GPT simply would not exist. It was open research and open sharing that got them to where they are today. Without those breakthroughs they couldn't have built their models, so when Amodei says they didn't get any help building their model, it's nonsense because their model is built on the back of countless open innovations.

Without open source AI many American companies would simple die off too. Take my small AI company, Kentauros, for example. Without open source models we'd be dead.

That's because we're building a product called AgentTutor that helps people collect the data they need to teach AI models how to do complex tasks on computers. We do automatic labeling of the data with foundation models. The terms of service of the large American AI companies would prevent us from doing this because they forbid using any outputs from their models for training other models.

Why? Again, because they don't want the competition. It's really as simple as that.

By having a big, powerful foundation model that can generate synthetic data to uplift newer, smarter models, they have the tremendous ability to hillclimb faster than everyone else while preventing you from doing the same and that prevents even good use cases like mine. I'm a part of little tech, not big tech. I'm not competing with OpenAI to build GPT. I couldn't even if I wanted to. Our models are designed to do one kind of agentic tasks really well.

Other use cases gets smashed by monopolistic licensing too. What if I was building a medical model to detect cancer and using GPT to do annotations of those X-Ray and MRI scans? I couldn't do it. I'd be locked out by the terms of service.

Open source changes this. It lets anyone innovate and dream up new ideas that can change the world.

At the Paris AI Action Summit, even former Google CEO, Eric Schmidt, has pivoted and called for the US to focus on open source or "China will ultimately become the open-source leader and the rest of the world will become closed-source."

The ugly truth is: this already the truth. China releases open models and all the US models are closed, except for Meta's models.

Every American University researcher and small AI business is dependent on Chinese open source models like DeepSeek R1, Qwen and Qwen VL and many more. That's because, other than Meta, American companies have abandoned open source. This is a looming disaster and it's one that some people not only don't want to avoid, they want to make it much, much worse.

And that takes us to the Jingoist.

Love It or Leave It

The jingoist has a very simple world view. Anything their country does is great. The jingoist is blind to any and all criticism about their own country. Their country is perfect and flawless always. It can do no wrong. Everything it does is always good and always the right thing, as long as it's "for the country."

This kind of thinking is behind the recent bill introduced by Republican Senator Josh Hawley, who's bill would criminalize use of DeepSeek and other Chinese models ***for Americans*** with jail time and fines of millions of dollars. This may be the single most idiotic bill introduced in a sea of AI related bills at the state and federal level over the last year and half, outpacing even the ill-fated California SB-1047 bill that was vetoed by the governor at the zero hour, despite pressing forward for months against massive opposition.

Unfortunately, he's not alone. The incoming Commerce Secretary, Howard Lutnick, said of DeepSeek: "Open platforms, Meta's open platform, let DeepSeek rely on it."

Leaving aside the fact that this is a completely wrong claim, because DeepSeek didn't rely on anything that Meta released at all, in any way, it's incredibly worrying that the incoming Commerce Secretary is this misinformed and doesn't seem to have even the most basic understanding of how open source works, why it works and why it was worth trillions to the world economy and millions of American jobs from coders to support.

Hawley's misguided and idiotic bill is music to the ears of the monopolist because they won't have to compete with open source models from China and they can jack up prices and muscle smaller American startups out of the picture, who rely on open source models because they didn't raise 10B to train a foundation model from scratch.

University researchers are particularly left behind without open source. Even prestigious universities like Stanford struggle to get the mega-computing resources they need for students learning the trade. The AI focused university that boasts industry AI Godmother, Fei Fei Li, as part of the its faculty, is just now spinning up a cluster of 248 H100s, a far cry from the 10s of thousands, 100s of thousands or even millions that US AI companies are working with now and in the coming years. Musk's xAI boasts 200,000 H100 and H200 chips.

In the US we only have one open-source champion, Meta. The bill would also put tremendous pressure on them *not* to release their models either, because China would get access to them too, which takes us back to our earlier realization that if Linux were born today it would have gotten strangled in the crib and never made it to becoming the bedrock of world software that it is today.

We need bills that do the *exact opposite* of this bill.

We need to let Americas make their own free choices about what models they use.

We also need to encourage more American companies to release open source and open weights models, not try to shut down the one that is releasing open weights models.

In other words, you counter influence *with* influence, not with restrictions.

Essentially, out of pure fear of China, the jingoist would kill off tens of thousands of smaller American startups, crush small business tech jobs and leave them with no alternative because they've criminalized open source.

(Breadlines in the Depression)

To counter the jingoist there is really only one argument that works: It must be framed as good for the country.

Luckily, open source *is* good for America. The jingoist, who looks at the world through a black and white filter of "my country good, other countries bad" can only be countered with the cold hard facts of what open means for their country. And here they are:

Open source creates trillions of dollars in value, powers American military and critical infrastructure and serves as the backbones of countless American companies big and small. It's used by as many as 99% of American fortune 500 companies. If AI is allowed

to develop properly and it doesn't get killed off by ludicrous legislation then it will serve as the backbone of the next American tech and business revolution.

Let them know that if open source is such a security risk, why does "the US Army [have] the single largest installed base for RedHat Linux" and why do many systems in the US Navy nuclear submarine fleet run on Linux, "including [many of] their sonar systems"? Why is it allowed to power our stock markets (NYSE, NASDAQ) and our clouds and our seven top supercomputers that run our most top secret workloads and the gov clouds that power all our of government agencies most important services like tax collection from the IRS which runs many services on AWS Gov Cloud?

It's likely that Senator Hawley doesn't know the sheer number of American companies and jobs that would get decimated by attacking open source and open weights AI and the last thing he or any other American congressperson wants to do is kill jobs. That's really the key for every American politician or any politician anywhere. Don't kill jobs. You want policies that create jobs and open source creates jobs.

So that's the way to fight the jingoist. It's maybe the simplest argument of all. Make it good for America or whatever country your local anti-open source politician hails from and you've got the only frame that's necessary to change their mind.

The thing is most people just do not understand and are not educated about what open source means to the modern world. They do not understand how critical it is to everything from our power grid to our national security systems and to our economies.

The main reason they don't understand it because because open source is *invisible*.

It runs in *the background*. It quietly does its job without anyone realizing that it's there. It just works. It's often not the interface to software, it's the *engine* of software, so it's under the hood but not often the hood itself. It runs our severs and routers and websites and machine learning systems. It's hidden just beneath the surface. It's an economic and engineering marvel of the modern world.

Open source is the most critical software infrastructure on the planet, bar none. That's why DARPA set out to understand "the most important software on Earth" in 2022. They wanted to understand it to make it even *more* secure and to understand why it works at all. “People are realizing now: wait a minute, literally everything we do is underpinned by Linux,” says Dave Aitel, a cybersecurity researcher and former NSA computer security scientist. “This is a core technology to our society."

It's not just a nice to have. It's absolutely critical to the functioning of the modern life at every level and to American security and the security of its allies. It's everywhere, an invisible bedrock beneath our entire digital ecosystem, underpinning all the applications we take for granted every day.

Show them that and you open their eyes to the light of the real world.

And so we come to the doomers, people who struggle to see the real world in any meaningful way.

Doomsday

Over the last few years battling the doomers, I've learned one thing:

You can't counter them directly.

You can't argue with them or convince them that AI will not kill everyone and take over the world like a megalomaniacal Marvel villain. Like Churchill said "a fanatic is one who can't change his mind and won't change the subject."

(Source: Marvel Studios, Ultron, robot villains are tired tropes in Hollywood)

Instead you've got to counter their influence, especially on legislators.

So far, we're mostly winning. The doomers were notably absent from the Paris AI Action Summit, when just two years ago it was called the AI Safety Summit. We fought back against doomer inspired bills like SB1047 in California and got it vetoed. Vice President Vance opened his keynote at the Summit by signaling a new era and a major policy shift:

> **"I'm not here to talk about AI safety, which was the name of the conference just a few years ago, I'm here to talk about opportunity."**

But the fight is not over. Despite their waning influence in the public and over legislator's minds, many doomer inspired bills are moving forward at the state level, as we saw earlier, and so the legacy of their all-out campaign to make AI fear front and center lives on like zombies in legislation that must still be mowed down.

I've written extensive counter-doomer arguments in my series of X posts called Destroying Terrible Anti-AI Arguments, Day One, Two, Three, Four, Five, Six. Don't waste them on doomers. Use them to help everyone else understand that they're a small, fringe group with a big megaphone.

In those posts I argued against, six major doomer fallacies:

1. "Would you open source a nuke?"
2. "What if AI released a deadly virus!"
3. "XYZ expert says AI is dangerous"
4. "My p(doom) ratio is X%"
5. "We need control frontier models like nuclear weapons with compute thresholds, model registration and inspectors."
6. "ASI will kill us all"

Here's the short version of one of the most common ones "would you open source a nuke"? — this is a classic case of begging the question. The question assumes AI is monstrous without bothering to prove it. It falsely compares a multi-use tool (AI) to a single-purpose weapon (a nuclear bomb).

AI is *not* a nuclear weapon. LLMs are not nuclear weapons. They're are not even close. A nuclear weapon is designed to do one thing and one thing only, kill as many people as possible. AI has a vast range of possibilities, from detecting cancer, to writing emails, to creating art, to doing tasks, powering robots that will clean the dishes and mop the floors, proving math theorems and more.

Very few technologies are inherently destructive like a nuclear weapon and so when you hear someone making that analogy you know they're making it from a lack of understanding or bad faith. They know it's not a real argument and they make it anyway because they have an ulterior motive which is usually to link AI to something horrific in your mind so they can justify restricting it and locking it down and censoring it.

A kitchen knife can cut vegetables or stab a loved one. 99.999% of people on the planet use it to cut vegetables and we don't hear cries to take kitchen knives off the market because one delusional nutcase stabs someone with it. We let the 99% cut vegetables and punish the person who stabbed someone.

AI is like that—a general-purpose tool that mirrors its user's intentions.

(Source: Swiss Army Knife from Victorinox)

Think of AI like Linux. As we saw earlier Linux powers supercomputers, smartphones, and routers, but it can also be used to write malware and launch DDOS attacks. Early critics, like Steve Balmer at Microsoft, tried to kill it by linking it to diseases like "cancer" or hated political ideologies like "communism." But we embraced Linux for its immense, widespread benefits anyways.

We don’t ban kitchen knives because a few use them to kill someone. We don't lock away every tool that might get misused. Instead, we mitigate risks through targeted, clear legislation that punishes the person who misused it, rather than criminalizing the tool itself.

I won't rehash all the anti-doomer arguments here. The key to understanding these counter-arguments is they're not to there convince the doomer. That's a waste of time. They're there to convince the lawmaker and the journalist and the everyday person who has never used ChatGPT that these terrifying black and white scenarios have nothing to do with actual reality.

Convince people seduced by doomer fear-mongering what the repercussions of those deluded belief systems will be. Let them know that if they let that fear take root, it will be a devastating loss for us all, because the revolutions in tech that made it so easy for people to get an education online, chat with our friends all around the globe, find information on any topic at the click of a button, stream movies, play games, find friendship and love are facing intense pressure by governments and activists the world over who want to clamp down on freedom and open access.

If the open source restrictionists win, we're entering an ugly world.

It's a world where your AI can't answer questions honestly because it's considered "harmful" (this kind of censorship always escalates because what's "harmful" is always defined by what people in power don't like), where information is gated instead of free, where open source models are killed off so university researchers can't work on medical segmentation and curing cancer (because budget conscious academics rely on open weights/open source models; they can fine tune them but can't afford to train their own) and where we have killer robots and drones but your personal AI is utterly hobbled and lobotomized.

Doomers are closet authoritarians. They believe in total and complete centralized control and that regular people like you and me can't be trusted to make our own decisions. If you're fighting their arguments you've got to let people know that it's openness and permissionless access that's behind the majority of the major technological, economic and societal revolutions of the last two hundreds years. Permissionless access. Permissionless innovation.

(Source: Soviet Propaganda poster of early pre-AI Doomer, Joseph Stalin)

What made the web so fantastic was an *open, decentralized approach* that let anyone setup a website and share it with the world and talk to anyone else *without* intermediaries. It let them share software and build on the software others gifted them to solve countless problems. It worked before beautifully. Why wouldn't it work again with AI?

Permissionless is the key driver of progress in the modern world and in history.

And it's the foundation of open source.

The basic idea behind open source is to give everyone the same building blocks, whether you're the government, a massive multinational corporation, a tiny individual academic researcher, a community sewing circle, or a small startup business.

When everyone has the same building blocks, its easier to build bigger and more powerful and more complex solutions on top of those blocks.

Take a technology like open source Wordpress. Wordpress made it super easy for anyone in the world to make and publish a beautiful website, even with limited design skills. It now powers over 43% of the web.

When I was younger it was hard to develop and publish a website. You had to do everything from scratch. But a wave of open technologies made it easier and easier as the years went by. We had the LAMP stack (Linux, Apache, MySql, PHP), an operating system, web server, database server and web scripting language, which gave many developers the common tools they needed to create more complex web apps. That, in turn, enabled the development of Wordpress on top of those technologies. Better and more powerful stacks followed, like the MERN stack, MongoDB, Express.js, React and Node, all open source technologies.

Each new layer of software makes it easier to build more advanced and more useful solutions on top of the last layer. It's the essence of progress and development. When open ecosystems are allowed to flourish and the enemies of open are finally beaten back once more, we'll see a flowering of new progress and innovations in a self-reinforcing virtuous loop just as we have in the past.

When you only have mud hut level technology you can only build one story tall mud huts. When you have a hammer and nails and standard board sizes you can build much taller and more robust structures. If you have steel and cement and rebar you can build skyscrapers stretching to the sky.

This same cycle of openness benefiting the world permeate the progress of the world again and again across many different aspects of life.

Take publishing:

When six publishers controlled publishing completely we had a flurry of the same kinds of books over and over again, all fitting a few basic models. Publishers kept 90% of the profits of writers. But the digital self-publishing revolution gave us a flood of amazing new

writers like Hugh Howey, with those authors now keeping 70% or more of their profits. Howey almost certainly would have never made it through the gate guardians of the big publishers to become a phenomenon with Wool and that was the fate of many authors before him. Gated access strangles progress.

Or take business and economics:

When starting a company required the permission of the king, you had a small group of super powerful corporations, like the British East India company that was once so mighty it conquered India with a private army that was *double* the size of the British army. That's right, it wasn't Britain that conquered India it was a giant mega corporation. And we think we have powerful companies now! Amazon ain't got nothing on the British East Indian company.

When breakthrough laws allowed anyone to create a company, we had a revolution of small businesses and a massive uptick in wealth and economic progress and longevity that has driven down poverty over two centuries to once unimaginable numbers and created more wealth than anyone in the 1700s could possibly even dream of and more.

Progress has a massive ripple effect that benefits the whole world.

Don't believe it? Just look at the world only a few hundred years ago.

Historian Michail Moatsos estimates that in 1820, just 200 years ago, almost 80% of the world lived in extreme poverty. That means people couldn't afford even the tiniest place to live or food that didn't leave them horribly malnourished. It means living on less than $1.90 a day in 2011 prices and $2.15 in 2017 prices. Actually you don't even need to go back that far. In the 1950s, half the world still lived in extreme poverty.

Today that number is 10%.

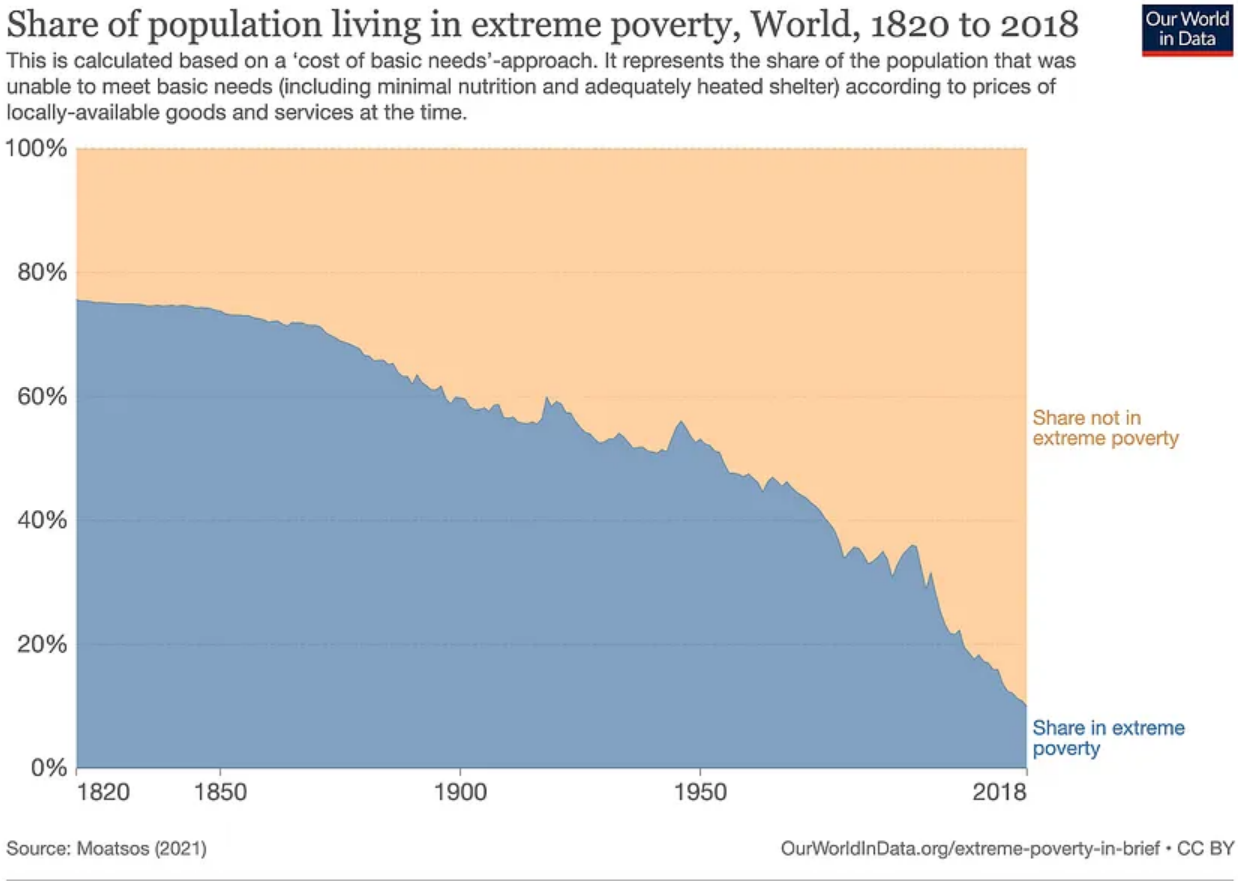
Share of population living in extreme poverty, World, 1820 to 2018
This is calculated based on a 'cost of basic needs'-approach. It represents the share of the population that was unable to meet basic needs (including minimal nutrition and adequately heated shelter) according to prices of locally-available goods and services at the time.
Our World in Data
100%
80%
60%
40%
20%
0%
1820
1850
1900
1950
2018
Share not in extreme poverty
Share in extreme poverty
Source: Moatsos (2021)
OurWorldInData.org/extreme-poverty-in-brief • CC BY

Nearly *half* of all children used to die, in every country on Earth, into the late 1800s, when economies were slower and more stagnant and less open.

In most first world countries infant mortality is now single digits or a fraction of 1% and even in the developing world it's 4%.

To see just how stark that drop is, take a look at this chart which shows child mortality over two millennia. You'll see that for most of human history it was a flat line of uninterrupted death and then it suddenly drops dramatically.

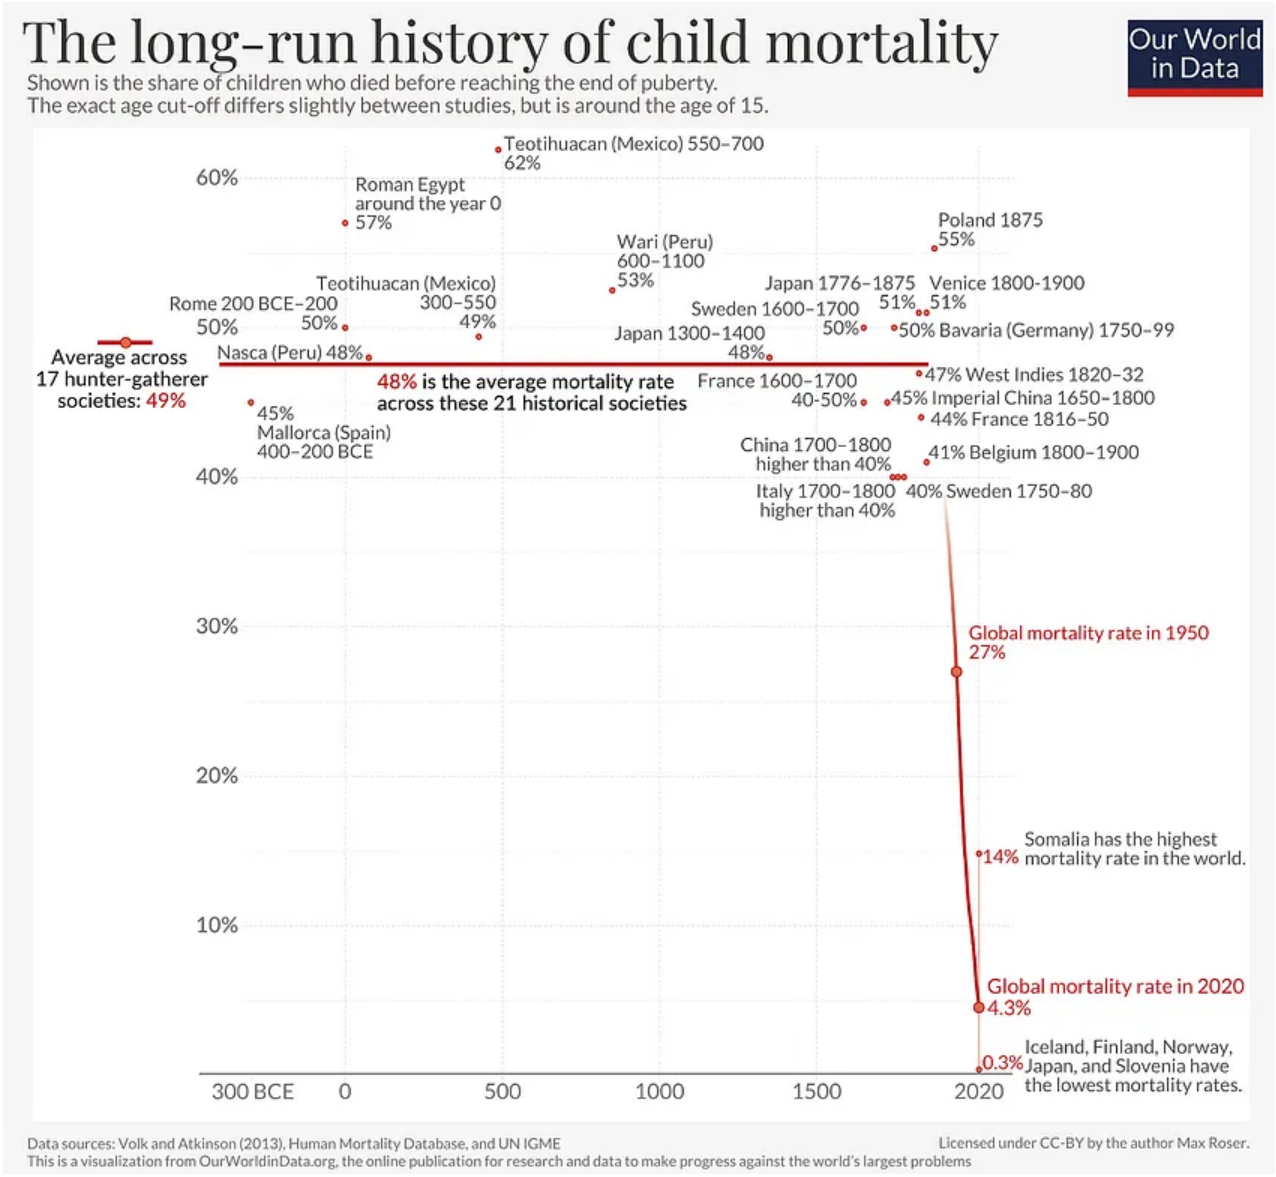

That's what progress does. It makes the world as a whole better. That's what openness, sharing and common building blocks do. They create possibilities for more progress and more ideas that build on the last ones. Once you have an understanding of microbes, you can build better defenses against them. Once you have microscopes, you can see those tiny menaces and understand if the drug you crafted to stop them is working. Each layer of knowledge and tools enables the next level of knowledge and tools. When anyone can build on the past and abstract those lessons to new domains we get new breakthroughs and new ideas and new businesses.

When you look back at the history of progress and openness and sharing and realize how crucial it is, it's easy to take it as a given that everyone can easily see just how important it is but they can't. The forces of closed and the enemies of innovation are a powerful chorus and they strike back again and again in history.

They believe in command and control. They believe its better to hoard knowledge and keep it secret. This is what doomers and other closet authoritarians believe and they get alone really well with the jingoists.

They image that if we just make something a little bit open, instead of all the way open, they we can prevent bad people from ever getting their hands on the best tools. But it never works that way.

That's because closed and gated systems have tremendous false positive and false negative rates. Gated systems are inexact, inefficient and they block many people you *want to have access* and *don't stop many of the people you don't want* to have access. Every time we try this in history we create the unwanted side effect of dramatically slowing down progress because information, ideas and software can't flow freely and get trapped between ugly, inefficient, governmental gates.

Crush doomers with facts and grounded, sane, clear policies and then watch their arguments vaporize like steam on a hot summer day.

And that takes us to the idiot.

The Idiot

The idiot is a special kind of enemy.

The problem is their arguments don't make sense on any level and we're wired to take an argument at face value, to assume there is at least *some merit* in what they're saying. But with the idiot there is nothing realistic or grounded in what they’re saying at all. There is no stable ground to stand on. It's just illogical nonsense from top to bottom.

If I'm shouting on the corner that the sky is made out of cheese and it's raining donuts, you don't need to try to logically pick apart the finer points of my argument because there are none.

Gary Marcus is one of the modern masters of idiotic arguments. He has a nearly perfect record of incorrect predictions and yet he's still regularly called upon by journalists and governments as an “expert” in AI, despite zero percent of the AI community acknowledging him as one.

Take this nonsensical Tweet he wrote the other day about the DeepSeek phenomenon:

> "Congress needs to bring in Zuckerberg and LeCun to discuss how their unilateral open-sourcing decision rapidly undermined the US advantage in Generative AI."

This is sheer idiocy because DeepSeek model wasn't based on any Meta models or any Meta ideas. In fact, the Chinese model outpaced Llama, Meta's flagship model, in just about every way. Not only did the DeepSeek team not need anything Meta released, the Chinese team created a number of their own breakthroughs independently, including a

novel, state of the art reinforcement learning algorithm, GRPO, that's on par with PPO (the current SOTA) but that uses a *lot less* memory and bandwidth and they trained a reasoning model, R1, that's on par with closed source OpenAI wonder model o1.

If you're looking for logic in what Marcus wrote, don't bother. But if you really want to follow the basic flow of his argument it's this: "Meta released open source; open source is bad; Chinese researchers must have copied what Meta did because Meta released open source; China only copies things; since I think open source is bad/dangerous, congress needs to call Meta to a hearing so that Meta can be stopped from releasing more open source because they're giving away our secrets."

In other words, neither the premise, the intermediatete reasoning, the conclusion, or the solution to the problem has *any* logic or merit *whatsoever* because they are all 100% wrong and based on nothing realistic in any way.

So how do you defend against an argument that has no basis in reality, doesn't reflect any known fact?

Unfortunately, the idiot is often the hardest person to argue with and their stupidity has a tendency to spread. Even a marginally intelligent person hears a stupid argument and says "well that's just stupid, no one will believe that" only to discover that the belief spreads like wildfire with other idiots. It seems self-obvious that what the person is saying is stupid and yet it's not obvious to many people and you often find idiots with massive social media followings. It's almost counter intuitive, the more stupid stuff they say, the more followers they get.

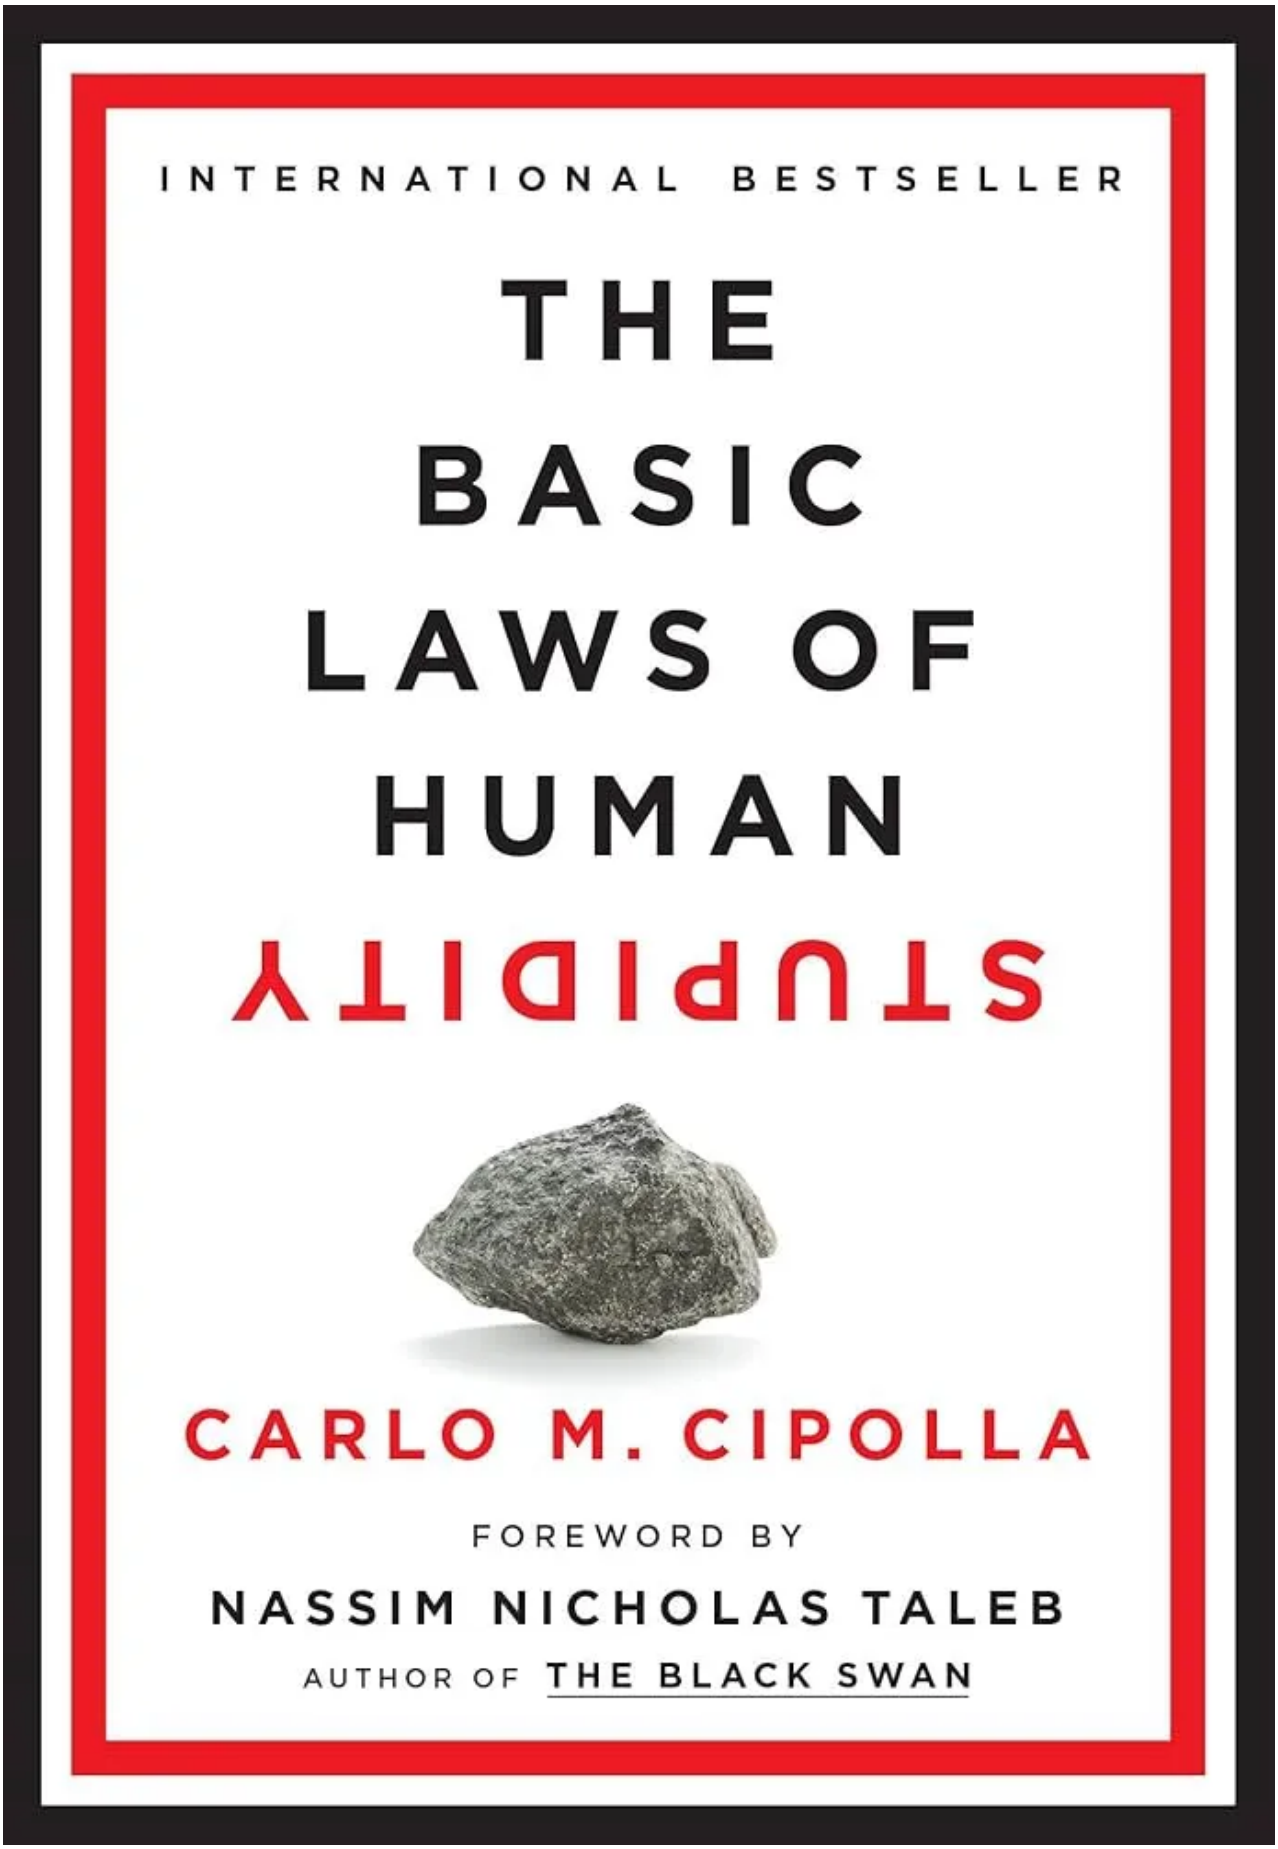
INTERNATIONAL BESTSELLER
THE BASIC LAWS OF HUMAN STUPIDITY
CARLO M. CIPOLLA
FOREWORD BY
NASSIM NICHOLAS TALEB
AUTHOR OF THE BLACK SWAN

The challenge of countering the stupid is put perfectly by Italian economist Carlo M. Cipolla in The Basic Laws of Human Stupidity:

> "[Smart and wise] people find it difficult to imagine and understand unreasonable behavior.
>
> "An intelligent person may understand the logic of a bandit. The bandit's actions follow a pattern of rationality: nasty rationality, if you like, but still rationality.
>
> "The bandit wants a plus on his account. Since he is not intelligent enough to devise ways of obtaining the plus as well as providing you with a plus, he will produce his plus by causing a minus to appear on your account. All this is bad, but it is rational and if you are rational you can predict it. You can foresee a bandit's actions, his nasty maneuvers, and ugly aspirations, and often can build up your defenses.
>
> "With a stupid person all this is absolutely impossible, as explained by the Third Basic Law. **A stupid creature will harass you for no reason, for no advantage, without any plan or scheme and at the most improbable times and places...**
>
> "...one cannot organize a rational defense, because the attack itself lacks any rational structure. The fact that the activity and movements of a stupid creature are absolutely erratic and irrational not only makes defense problematic but it also makes any counterattack extremely difficult—like trying to shoot at an object that is capable of the most improbable and unimaginable movements."

Cipolla, Carlo M.. The Basic Laws of Human Stupidity (pp. 53-54). (Function). Kindle Edition.

In other words, the very definition of stupidity is someone who causes you a loss and potentially even themselves a loss for absolutely no reason at all.

So how do we fight that?

The best and only argument against the idiot's attacks on open source AI or anything else, is to simply dismiss them as absurd and let people know that just because the person might have credentials or a long resume their argument is pure and utter nonsensical crap that doesn't even deserve a response because it's like arguing about whether the moon is made out of cheese.

Simply say there is nothing to see here and move on. Pretend they don't exist. Don't give them any attention at all and counter their influence elsewhere.

Unfortunately, this strategy doesn't always work. And the reason is that sometimes the idiotic argument comes from a high status individual. That's someone who's otherwise credentialed or smart or revered for their accomplishments in other way.

Yoshua Bengio is masterful at making stupid arguments about the future of AI *and he is also* one of the godfathers of machine learning and his contributions to the field of machine learning can't be underestimated. Sadly, he recently fell under the spell of the doomers and now spends much of his time campaigning about how we might "lose control" of AIs or they might "take over."

With a high status person you can't just dismiss them as outright idiots like you can with charlatans who've contributed nothing to field. They must be taken more seriously. But there's another rule of the Basic Laws of Human Stupidity comes into play:

"Stupid is independent of any other characteristic in a person."

In other words, someone can be a Noble prize winner or a fry cook and still deliver an utterly idiotic argument.

Take this recent tweet from Bengio:

Yoshua Bengio
@Yoshua_Bengio
While the AI Action Summit was the scene of important discussions, notably about innovations in health and environment, these promises will only materialize if we address with realism the urgent question of the risks associated with the rapid development of frontier models.
Science shows that AI poses major risks in a time horizon that requires world leaders to take them much more seriously. The Summit missed this opportunity.

Here he argues that "Science shows that AI poses major risks." Science does nothing of the sort. Maybe science fiction does but that's not anything to base policy on. Stories are about one thing only, conflict. They don't reflect real life, only *a single aspect of life* and that makes them a terrible lens to see the world. You're guaranteed to see life through a darker lens if you view it only through stories.

Science, on the other hand, is a way of getting closer to empirical truth by having multiple people across the world independently verify an observation. There is no science that accurately predicts the future of how society and technology develops.

Science might verify the chemical structure of a newly discovered material or verify the efficacy of a drug to combat a disease. But there is no science that shows that AI will rise up and take over the world or kill us all or anything of the like. There are rampant and wild

speculations but none of this is based on any firm evidence of any kind. So Bengio is either being deliberately disingenuous here in order to lend the gravitas of science to his belief system about evil AI going rogue or he's just making an idiotic argument.

Here's another one from a snippet of one of his speeches:

> He says that once OpenAI has a truly powerful model they won't release it and instead they'll use it to "basically wipe out other companies and other economies that don't have these [models.]"

This is another incredibly terrible argument. It starts off with a reasonable premise, that an AI company might be economically tempted to keep their most powerful models to themselves and not sell direct access to them. That makes sense. It's a real possibility. This was the case with Google for many years, where AI powered their products like search but users never got direct access to it.

But after that the argument leaps into the absurd and suddenly OpenAI is ravaging whole companies and economies with no intermediate steps. How did we go from 1 to a million in a single leap?

In other words, in his scenario nobody else on Earth has created a powerful AI, only OpenAI in a vacuum. Also no other companies can defend themselves with their own adaptations, there are no legal or societal pushbacks or mitigations and entire economies go under. This is absurd. There are always counter balancing forces to life. As we saw with DeepSeek R1, even when OpenAI tries to conceal their most powerful training breakthroughs, other teams can replicate it in a few months through reverse engineering and logical deduction.

In other words, nothing happens in a vacuum. There will be pushback and friction and other companies and people with breakthroughs in AI at the same time so this argument by Bengio is another swerve into the absurd.

Unfortunately, the problem with a high status person like Bengio making this arguments is people are prone to thinking that expertise in one area applies equally to other independent areas.

They do not.

Just because someone may have forgotten more about machine learning than most people will ever know does not mean they're the least bit qualified to predict how technology effects complex systems like human society over time.

Whenever you hear something like “so and so is expert in AI and they say AI will take over the world,’ the first question to ask yourself here is "expert in what? Just because you are an expert in the inner workings of AI, does not mean you are an expert in psychology, business, economics, history, or how technology develops and changes over

time. That's generally the work of sociologists, futurists and historians (who look for patterns in the past and sometimes try to apply them to the future). And even then it's a very inexact discipline.

So the answer to defending open source against the high status individual's idiotic argument is to acknowledge that person's status and contributions to the world but to crush the argument itself and remind people that just because you know something about one thing doesn't mean you get automatic credit in another area.

And if the person is just a general charlatan and grifter, with no real status whatsoever, just ignore them and move on to counter their garbage with other people so their idiocy doesn't spread.

You Are the Revolution

Microsoft once tried to kill open source. They launched a massive anti-Linux marketing barrage and tried to link it to disease and communism. They failed and now open source is the most successful software in history, the foundation of 90% of the planet's software, found in 95% of all enterprises on the planet.

Imagine if they'd been successful in their early attacks? They would have smashed their own future revenue through short-sightedness and total lack of vision.

Yet here we are again, with a witches cauldron of people who want to smash open source and open weights AI out of ignorance, blind animal panic and fear:

- Ultranationlists who think that if China gets to use something open too then it must be stopped at all costs, including smashing our own businesses and progress
- Aging AI safety fetishists like Max Tegmark, who worry AI will get "out of control" who want to crush and control AI development at every turn because they fear the end of the world like some bad sci-fi novel
- American AI lab heads like Dario Amodei of Anthropic who would rather strangle the competition and create monopolies than compete fairly on an even playing field

They must be stopped. It's really as simple as that.

If we let monopolists and the delusional drive policy in the world, nothing good ever comes from it.

If you want a world where small AI businesses are locked out of competition with the big AI labs, where Chinese models are banned in America, the land of freedom and opportunity, where your models are censored and controlled and what you can and can't do with them is strictly controlled, then embrace the jingoists, doomers and monopolists.

But if you want the next multi-trillion dollar revolution in software to benefit the whole world, you've got to fight them at every step until they're stopped for good and the technology can develop naturally and organically.

Time has a way of handling these problems. Given enough time, earlier idiotic laws will get rolled back and fanciful arguments about the end of the world will fall away as people experience AI in the real world. R1 and o1 were released months ago and the sky has not fallen. When people's experience don't match up with the doomsday predictions they'll adapt easily.

(Source: Midjourney, No Chicken Little the Sky is NOT falling)

AI will just be another thing to tomorrow's children, like a tree or a rock or a smart phone. It will just be there and they'll use it and understand it. If you tell them that some people once thought it would kill us all they'll just look at you like your nuts and wonder what the hell you're talking about.

Maybe you think that lots of people have tried ChatGPT or another advanced model but the reality is that most people in the world just haven't. As more and more regular people get their hands on AI and see its benefits and realize the sky isn't falling, and that no machines have risen up to take over and the world hasn't ended, public opinion will naturally shift over time away from crazy fear and insanity.

We also need more companies to get back to embracing open source.

Fantastic models like Allen Institutes Molmo and Olmo. Amazing Chinese models like Qwen and Deepseek. American models like Llama.

We need *a lot more* American companies to embrace open source, not fight it because they want to lock up control of their models for a decade. And we need new countries getting into the picture. German models. Indian models. South African models. Saudi Arabian models. Italian models. Models from everywhere. We need a common intelligence infrastructure that we can benefit from the same way we benefit from open source software today.

AI will deliver an intelligence layer for the world.

But we might never get there if the jingoists and the stupid and the monopolists have their way with us.

Today's critics of open source AI trot out the same tired old tropes, like "security" or "China = bad" and they tell us only a small group of companies can protect us from our enemies so we've got to close everything down again and lock it all up behind closed doors.

Don't believe them.

Remember that the ultimate definition of stupid is people who do damage to others without any benefit to themselves and maybe even a net loss to themselves.

- The jingoist would kill Linux so it never runs on Chinese supercomputers, ensuring that it never powers their own super computers as well, aka cutting off their own nose to spite their face.
- The doomer would kill AI because their fear based creatures who read too many sci-fi books and because they want to feed their own egos and feel like they're the noble hero saving the world from the apocalypse.
- The monopolist just wants to make sure they get all your money and don't have to worry about anyone else competing with them on a fair and open playing field.

- The idiot just wants to cause damage to themselves and everyone else for no gain at all.

If the only people working on AI have war chests of 10B or 100B you're going to have a very sad AI ecosystem, composed of a few powerful players with very limited and narrow thinking. That's not a world that we want. That's not America. America is about robust competition and the fact that any person regardless of where they came from can make a go at it and make it big.

We want competition and we want more minds working on intelligence. We want an intelligence revolution. We want ambient intelligence everywhere, embedded into every layer of the world economy, making it faster, better and more efficient. We want AI that rapidly creates new drugs that cure diseases and solves centuries old problems.

And there's only one way to get there.

Open.

Future History is a reader-supported publication. To receive new posts and support my work, consider becoming a free or paid subscriber.

Share

Discussion about this post

1 more comment...

No posts